



Enumerating Unnumbered

Saving IPv4 addresses
because it's 2023

whoami

- FRR (...Quagga...) maintainer since the iron age
- working for NetDEF since ≈2012
- non-profit in the middle of FRRouting (CI, RelEng/pkggs, review, feature contracts)
- find Martin Winter @ RIPE OSS WG



Outline

1/3: explicit PtP addresses

2/3: intermediate hacks

2½/3: IS-IS & BGP intermezzo

3/3: IPv4 with IPv6 nexthops (RFC 5549/8950)

4/3: General caveats



Note

configuration references are for FRRouting but similar options will exist in other vendors

all of this should be available for (reasonable) vendors (and interop)

ECMP is sometimes broken, test if needed



Unnumbered 0/3: the stone age

- classful IP networks \Rightarrow nothing special for PtP links \Rightarrow reuse address from another link

```
interface FastEthernet0/0/0  
  ip address 192.0.2.34/24
```

```
interface Serial0/0/0  
  ip unnumbered FastEthernet0/0/0
```



Unnumbered 1/3: PtP addresses

- configure explicit local and remote IPv4 addr
- (almost?) all platforms these days can do this on LAN links

```
interface Loopback
```

```
  ip address 192.0.2.34/32
```

```
interface Ethernet0/0/0
```

```
  ip address 192.0.2.34 peer 192.0.2.56/32
```



Unnumbered 1/3: PtP addresses

- pro: should be available & interop everywhere
- con: neighbor's address is in the config
...and sets a connected route that may break
- con: ECMP may be broken if vendor didn't consider it/test it (especially with BGP)



Quick warning

192.0.2.34 peer 192.0.2.35/32
is not the same as
192.0.2.34/31

```
> Internet Protocol Version 4, Src: 10.255.0.1,
  ✓ Open Shortest Path First
    > OSPF Header
    ✓ OSPF Hello Packet
      Network Mask: 0.0.0.0
      Hello Interval [sec]: 1
```

```
> Internet Protocol Version 4, Src: 10.199.1.0,
  ✓ Open Shortest Path First
    > OSPF Header
    ✓ OSPF Hello Packet
      Network Mask: 255.255.255.254
      Hello Interval [sec]: 1
```



2a/3: lone /32 on the iface

2b/3: special config option

```
interface Loopback  
  ip address 192.0.2.34/32
```



```
interface Ethernet0/0/0  
  ip address 192.0.2.34/32
```



```
interface Ethernet0/0/1  
  ip unnumbered
```

} depending on the vendor,
only one of these may work
(or they do the same thing)



2ab/3: same as PtP really

- OSPFv2 has no special “unnumbered” mode
- outcome should in fact be identical

```
> Internet Protocol Version 4, Src: 10.255.0.1,
  ▾ Open Shortest Path First
    > OSPF Header
    ▾ OSPF Hello Packet
      Network Mask: 0.0.0.0
      Hello Interval [sec]: 1
```

```
> Internet Protocol Version 4, Src: 10.255.0.1,
  ▾ Open Shortest Path First
    > OSPF Header
    ▾ OSPF Hello Packet
      Network Mask: 0.0.0.0
      Hello Interval [sec]: 1
```



2ab/3: same but better

- peer's IPv4 is not in your config
- no possibly-broken connected route for peer

only problem: can't (easily) run BGP anymore



2¹/₂/3: IS-IS and BGP

- IS-IS just works™
(did anyone expect anything else?)
- BGP... needs some way to configure a session
⇒ works with PtP but not “full” unnumbered connected route “problem” is also “solution”



3/3: BGP invents a better(?) wheel

- let's just use IPv6 nexthops for IPv4 routes
⇒ RFC5549/8950 (Extended Next-Hop Encoding)

```
# ip route list
192.0.2.0/24 nhid 1 via inet6 fe80::1234 dev dummy0
# ip nexthop list
id 1 via fe80::1234 dev dummy0 scope link
```



3/3: The BGP wheel has corners

- Autodiscovery is a separate, vendor-specific feature (RA, ndisc or LLDP based)
- fate sharing is broken — use BFD if possible
- theoretically possible to run router without any IPv4 address... *theoretically*¹

¹ Your mileage may vary depending on your driving style, phase of the moon, color of your 19" rack, and keyboard layout used while configuring your router. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES. Why are you still reading this. It should be clear by now that this is a joke. Please give your router at least one IPv4 address on loopback.



3/3: BGP 5549/8950 config

- generally needs interface and some “enable” switch, sometimes also peer IPv6 address
- more available-ish on datacenter devices

```
router bgp 64500
  neighbor fe80::1234 remote-as 64501
  neighbor fe80::1234 interface vm123
  neighbor fe80::1234 capability extended-nexthop
```



Quick warning, round 2

- Some vendors only support this with link-local addresses.
- Trying to use GUA or ULA IPv6 for this is the “wrong thing”™ anyway.
- If you remove all IPv6 LL addrs, you are breaking your network. Stop doing that.




```
kvm3.netdef.org# show ip bgp
```

```
... ✂ ... ✂ ...
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 0.0.0.0/0	0.0.0.0	0		32768	i
*> 194.147.139.17/32	fe80::5054:ff:fg h i:jk l m	0		0	65001 i
*> 194.147.139.18/32	fe80::5054:ff:fg h j:km n p	0		0	65002 i
*> 194.147.139.22/32	fe80::5054:ff:fg i k:nq r t	0		0	65003 i



3/3: It's super effective

- Just to point out: this is a great way to route individual /32s into VMs (and provide a default route the other way.)
- the cool kids are doing this already
- also: <https://github.com/euro-ix/rfc8950-ixp>



3³/₄/4: summary

OSPF

IS-IS

BGP

explicit PtP



“unnumbered”



IPv6 nexthop



4/3: General caveats

- source address selection & traceroute
- sshing to the other end for debugging
- anything that identifies links by IPv4 address



4/3: source addr sel. & traceroute

- search order is: outgoing interface, loopback, other interfaces
- make sure you have a “good” IPv4 on loopback
- if the router has “weird” things, may need to add /32 on interfaces (same addr everywhere)



4/3: Debugging

- Use the fe80, Luke
- It's one hop only ... *but* ...
very hard to break / lock yourself out,
regardless of unnumbered setup.



4/3: Link identification

- RSVP-TE and LDP sometimes use IPv4 addresses to identify a link
- some SDN controllers too
- no way around it, need to fix or get rid of those 😞



Questions?

